LegUp: Accelerating Memcached on Cloud FPGAs

Xilinx Developer Forum
December 10, 2018

Andrew Canis & Ruolong Lian
LegUp Computing Inc.

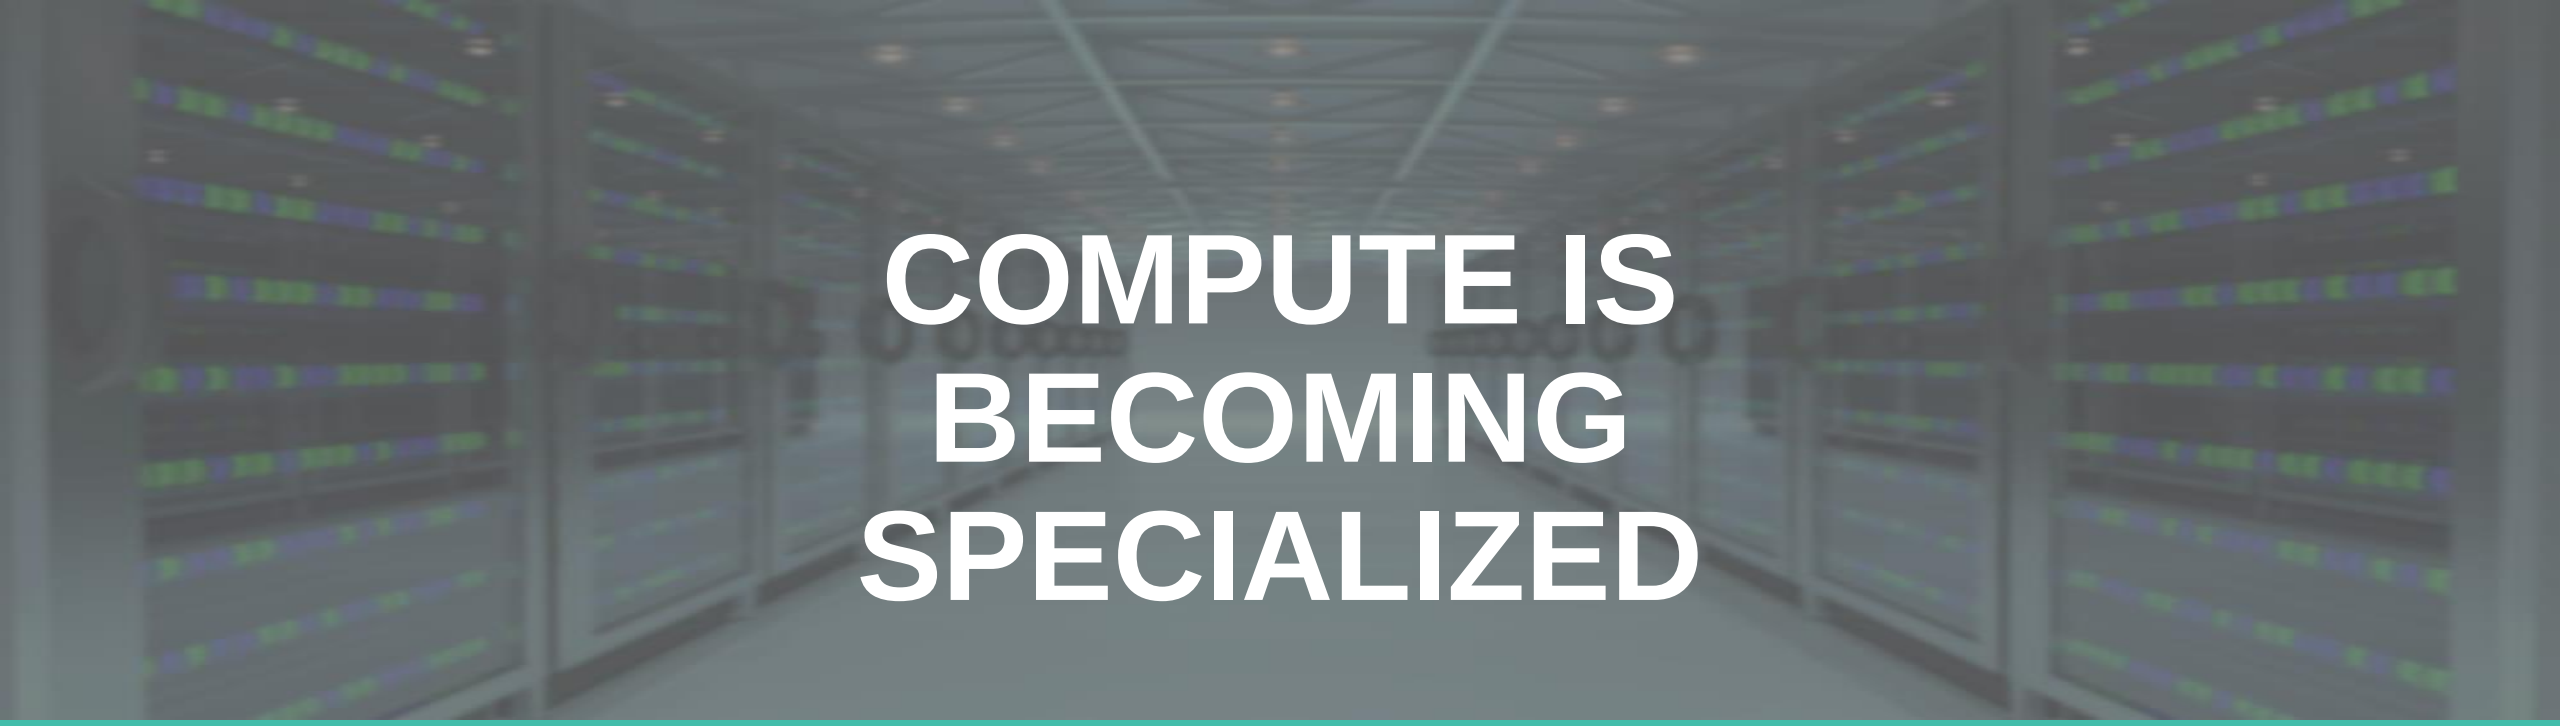

COMPUTE IS BECOMING SPECIALIZED

GPU

Nvidia graphics cards are being used for floating point computations

TPU

Google tensor processing unit used for machine learning

FPGA

Reconfigurable hardware.

FPGAs excel at real-time data processing.

LEGUP HLS PLATFORM

Software

Software
Test/Debug

LegUp

A Unified Hardware Acceleration Platform

Hardware
System

Hardware

Vendor
Agnostic

CPU

XILINX
ALL PROGRAMMABLE

SDAccel
Environment

LegUp

The Era of FPGA Cloud Computing is Here

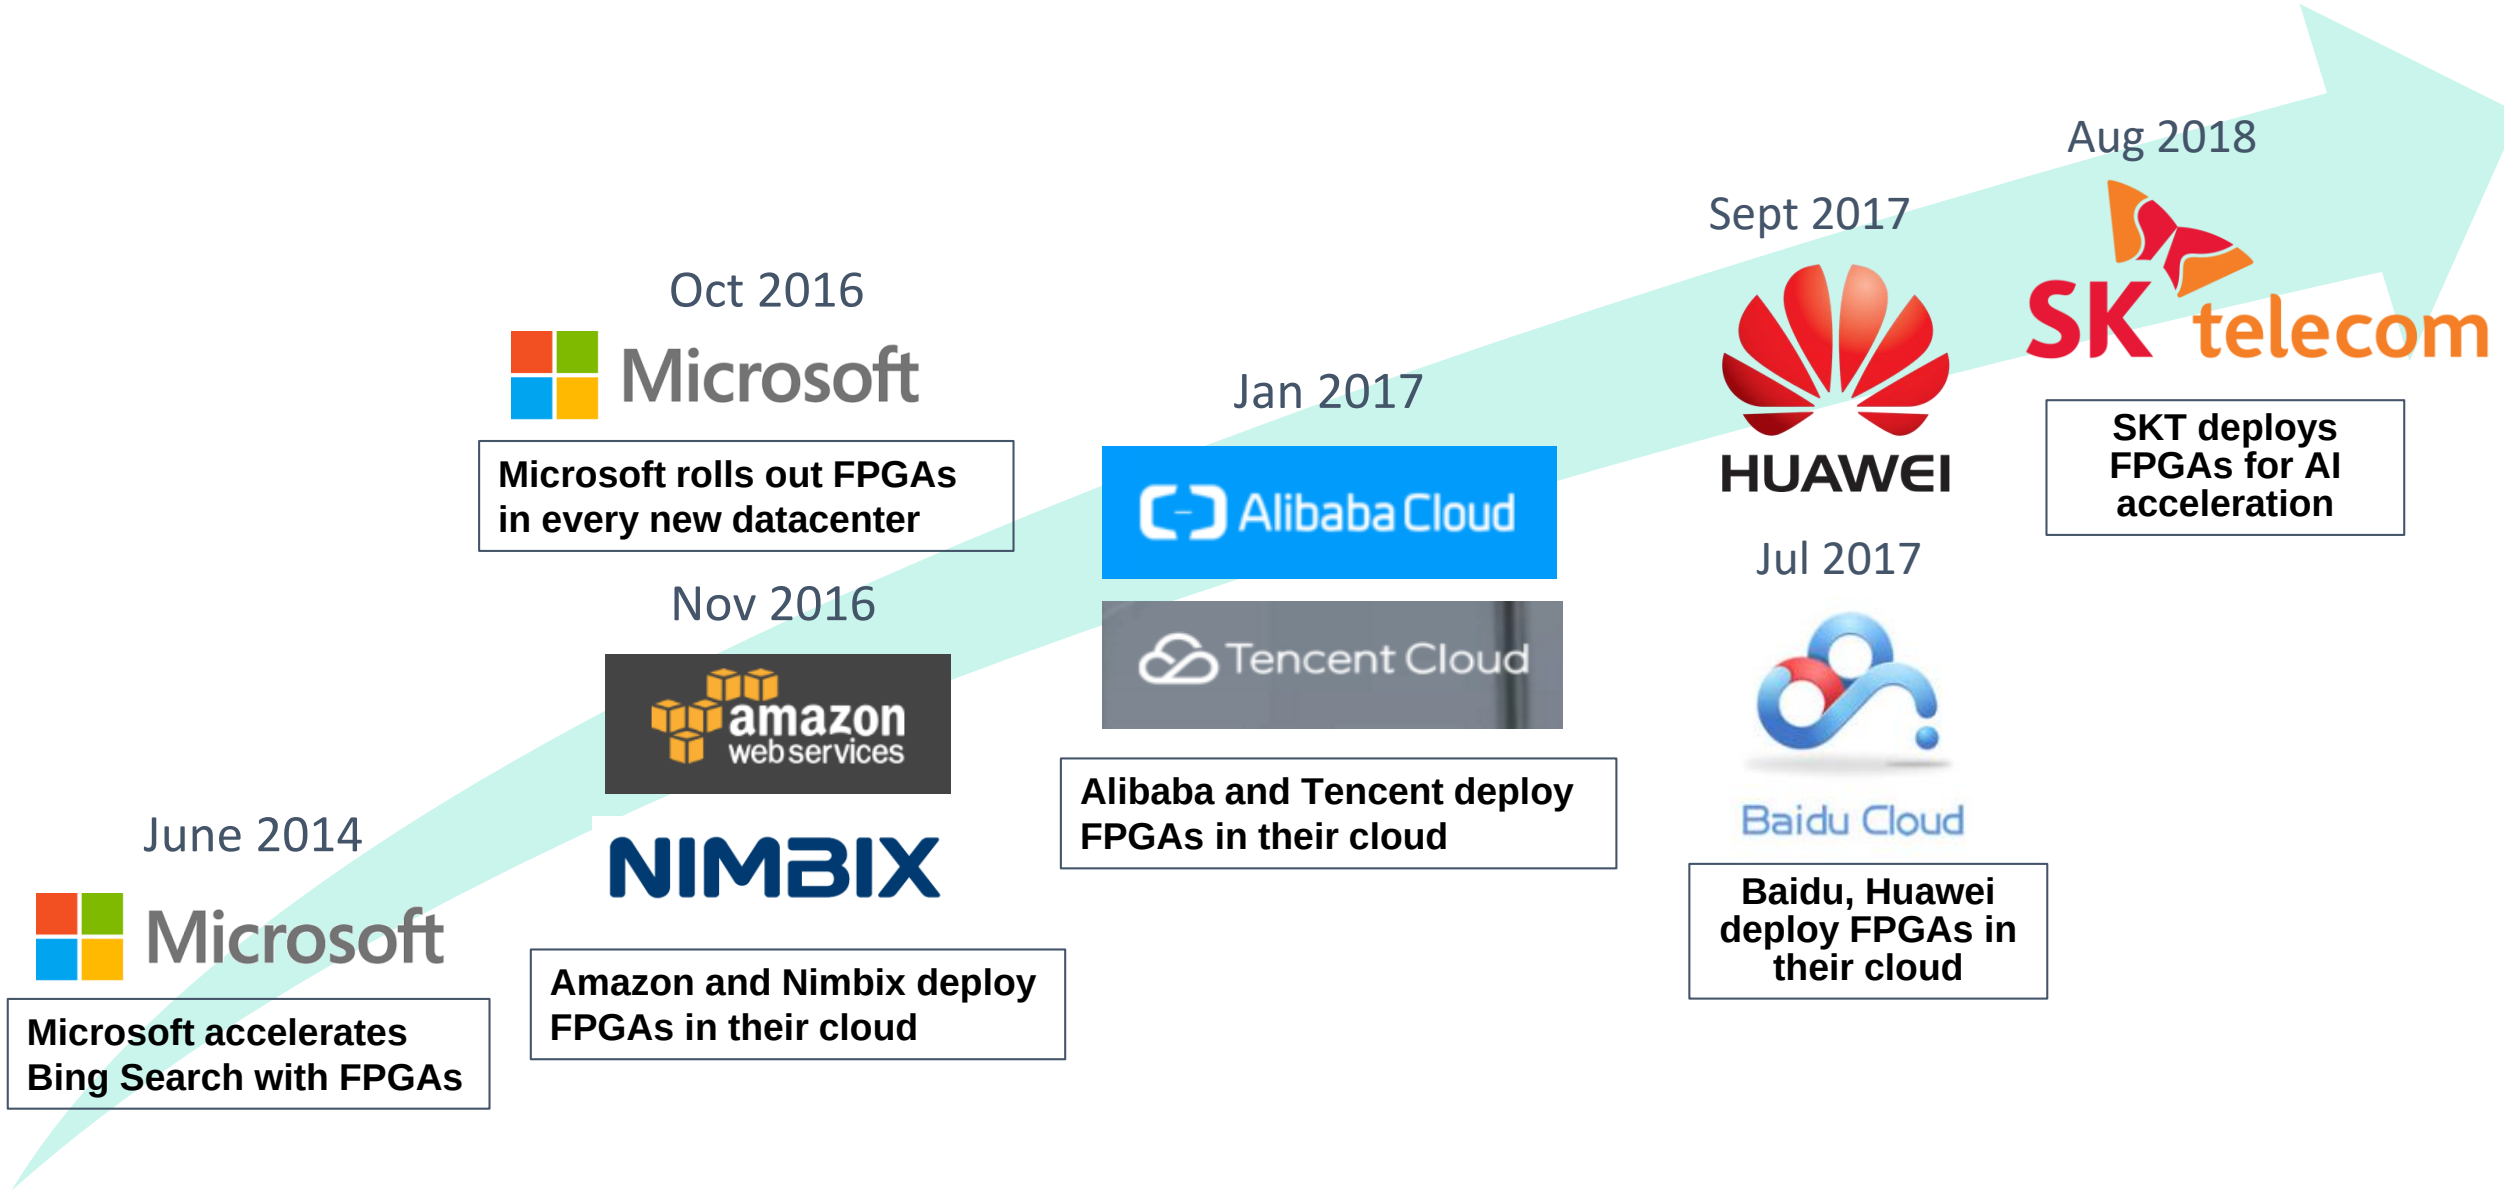

CLOUD PLATFORM

- Network processing engines on cloud FPGAs and on-premises FPGA acceleration cards

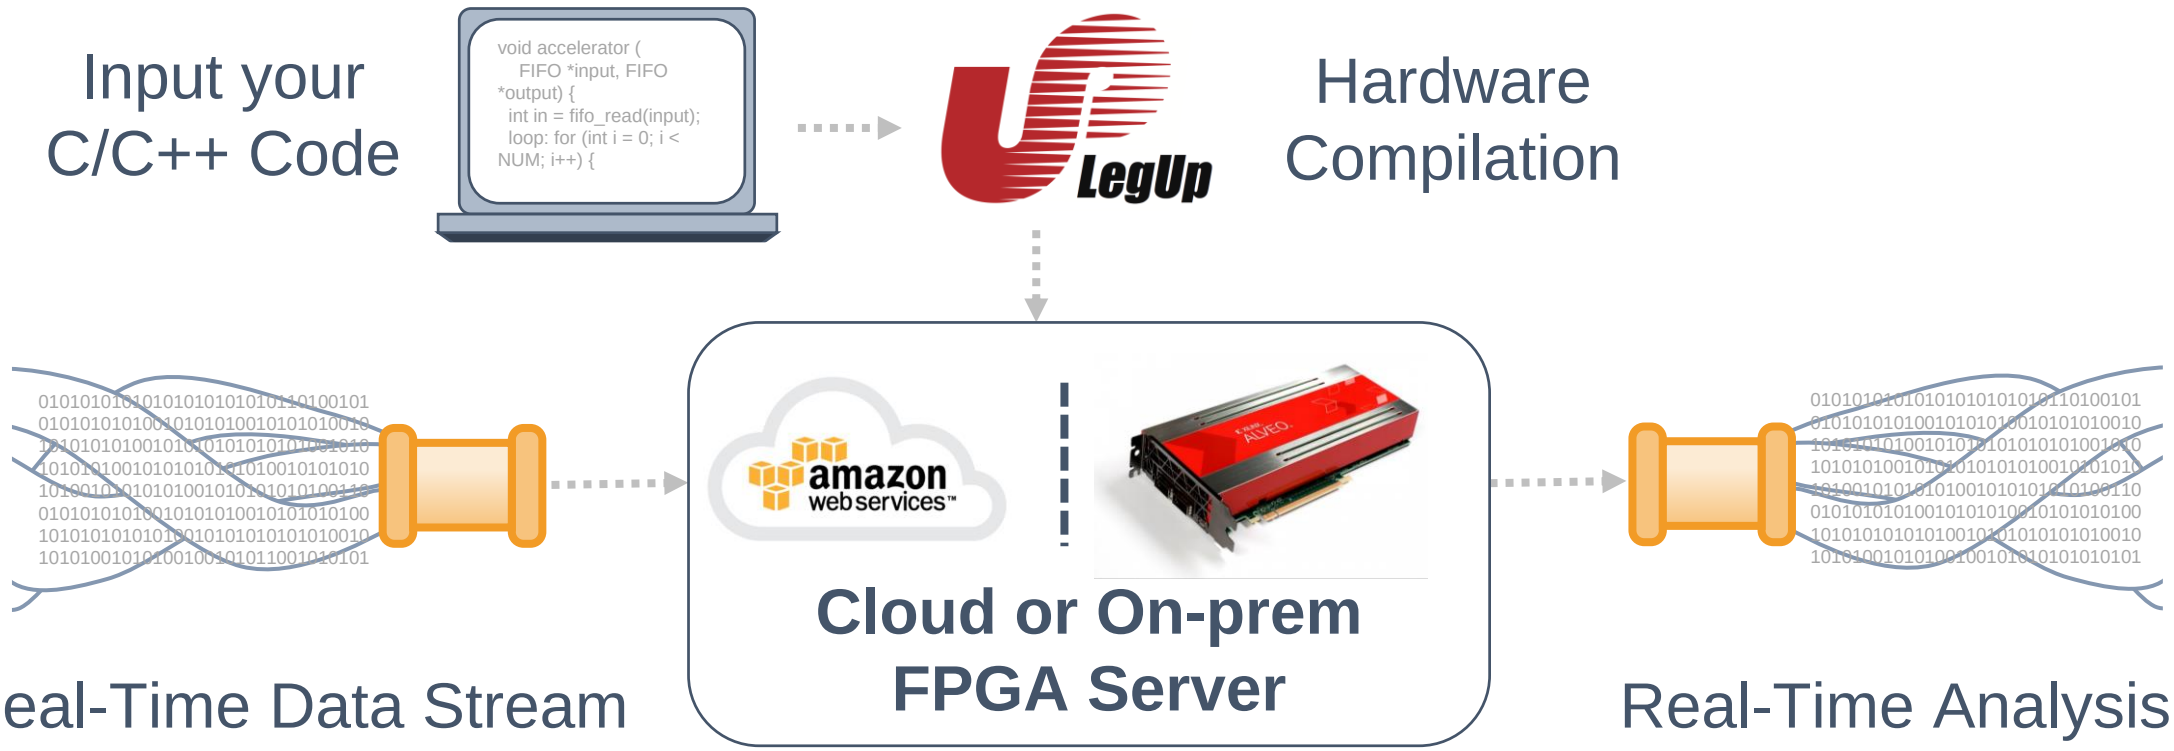

What is Memcached?

- Memcached is a distributed in-memory key-value store
  - Used as a cache by Facebook, Twitter, Reddit, Youtube, etc
  - Facebook Memcached cluster handles billions of requests per second
- Memcached Commands:
  - Set key value
  - Get key
- Typical deployments:
  - Amazon ElastiCache
  - Google Cloud App Engine
  - Self-hosted
- Easy horizontal scaling:
  - Cluster of Memcached servers handles the load

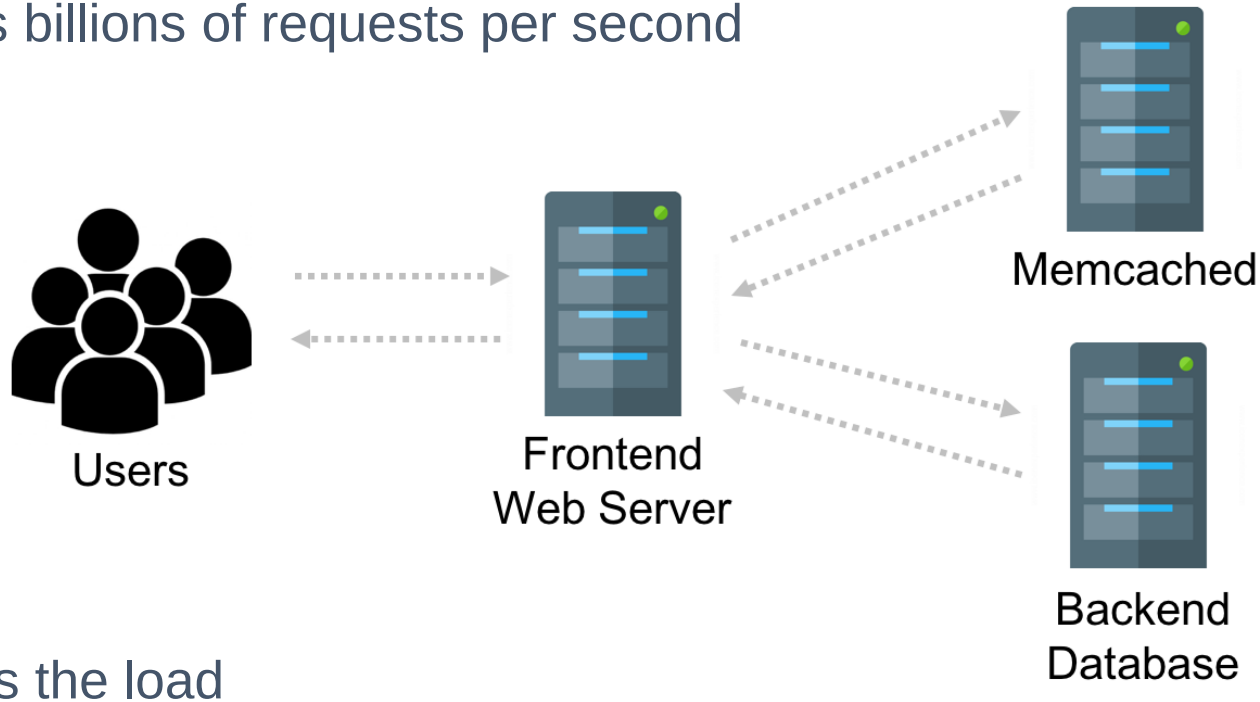

Introducing: World's Fastest Cloud-Hosted Memcached

- Easy to Deploy
- Lower TCO
- 10Gbps network

amazon web services™

Powered by AWS FPGAs with LegUp's Platform

9X
HIGHER REQUESTS/SEC

9X
LOWER LATENCY

10X
LOWER TCO

Memcached vs. AWS ElastiCache

- Benchmarked Memcached against AWS ElastiCache
  - AWS provides a fully-managed CPU Memcached service
  - Different instance types based on RAM size, network bandwidth, and hourly cost
  - Chose an ElastiCache instance with the closest specs to F1

AWS Instance	vCPUs	RAM	Network Speed	Cost
cache.r4.4xlarge (CPU)	16	101 GB	Up to 10 Gbps	**$1.82/hour**
f1.2xlarge (FPGA)	8	122 GB	Up to 10 Gbps	**$1.65/hour**

Experimental Setup

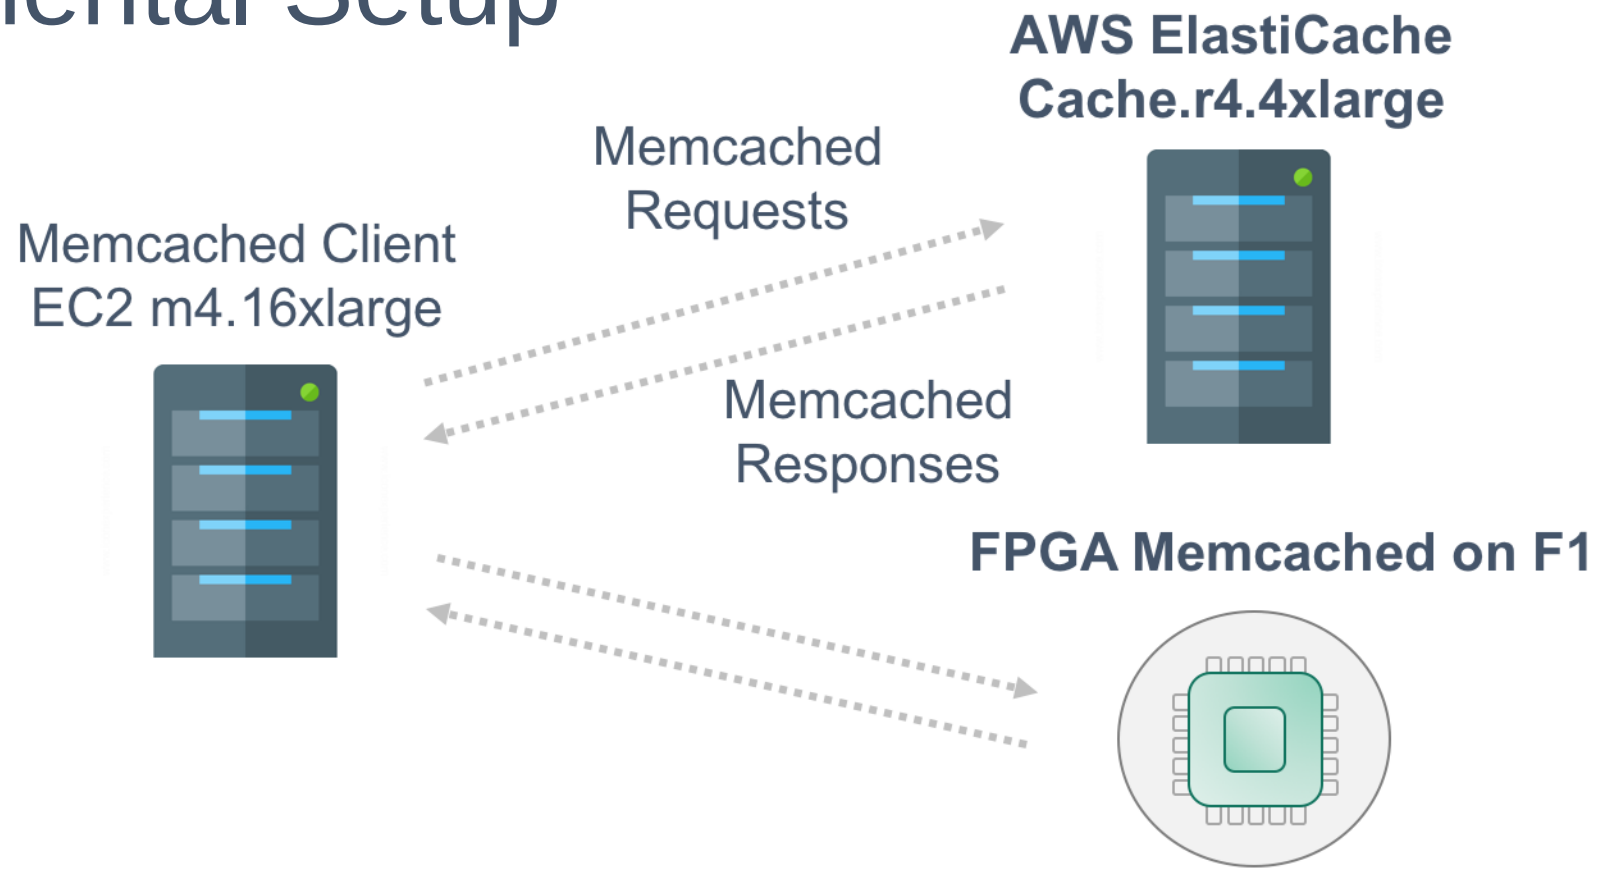

- Memtier_benchmark: Open-source Memcached benchmarking tool
- 100-byte size data, pipelining (batching) of 16
- Varied number of connections to Memcached

LegUp

Throughput Results

- Up to **9X** better ops/sec vs. ElastiCache

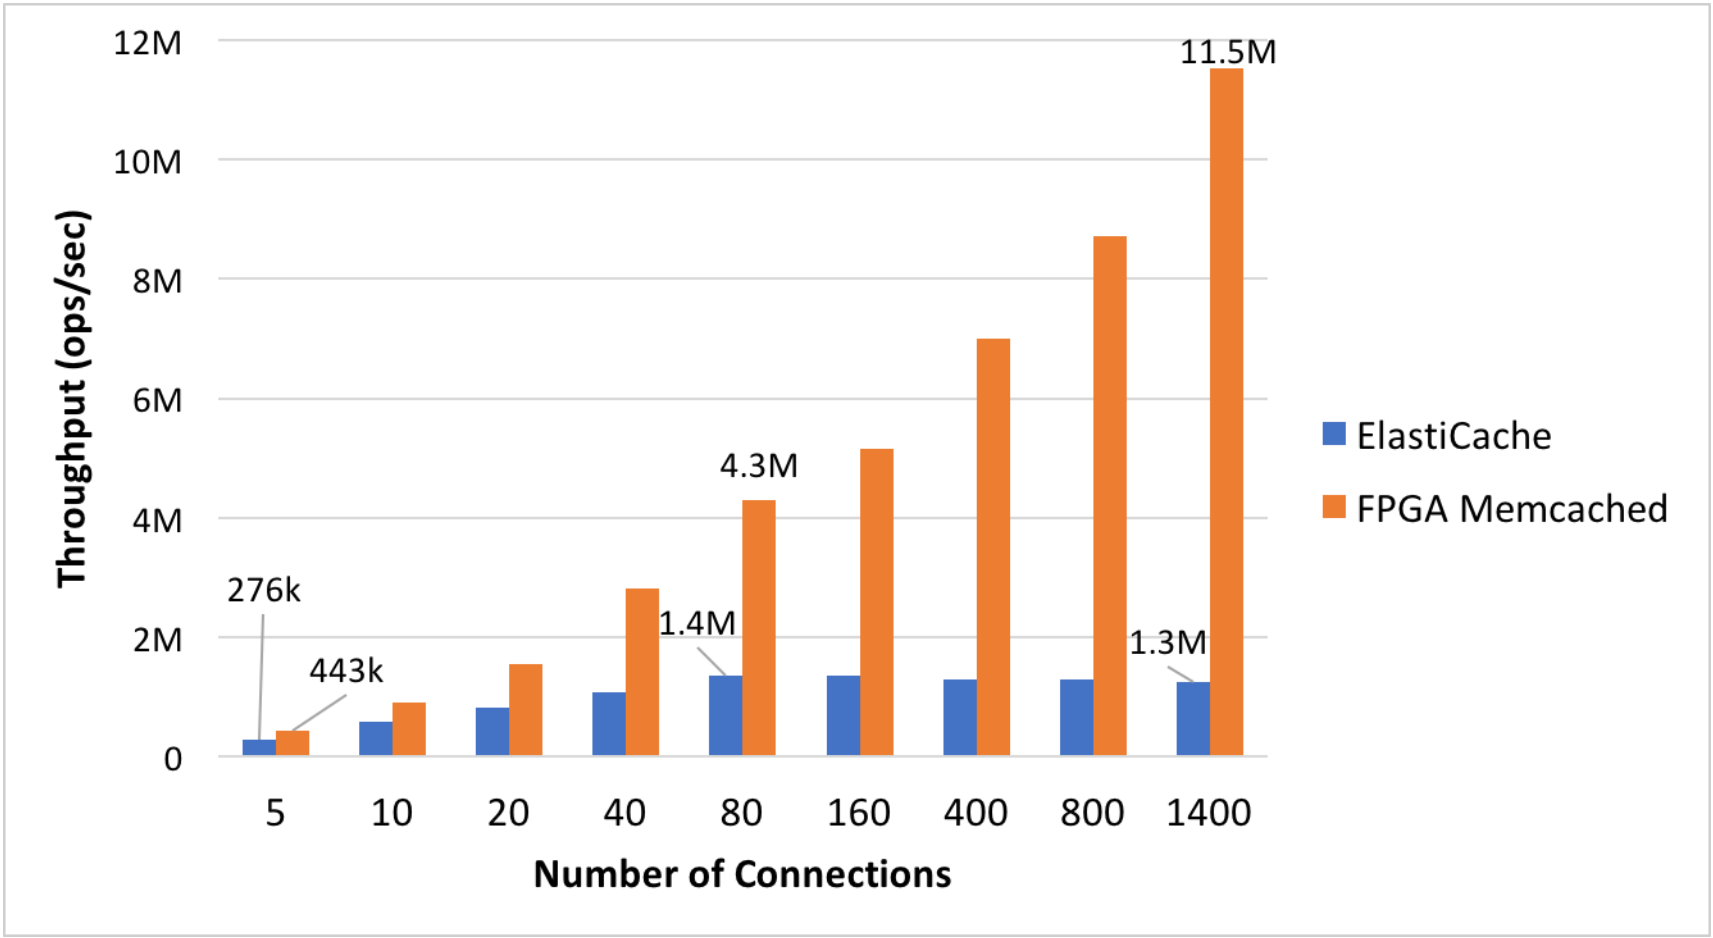

LegUp

Latency Results

- Up to **9X** lower latency vs. ElastiCache

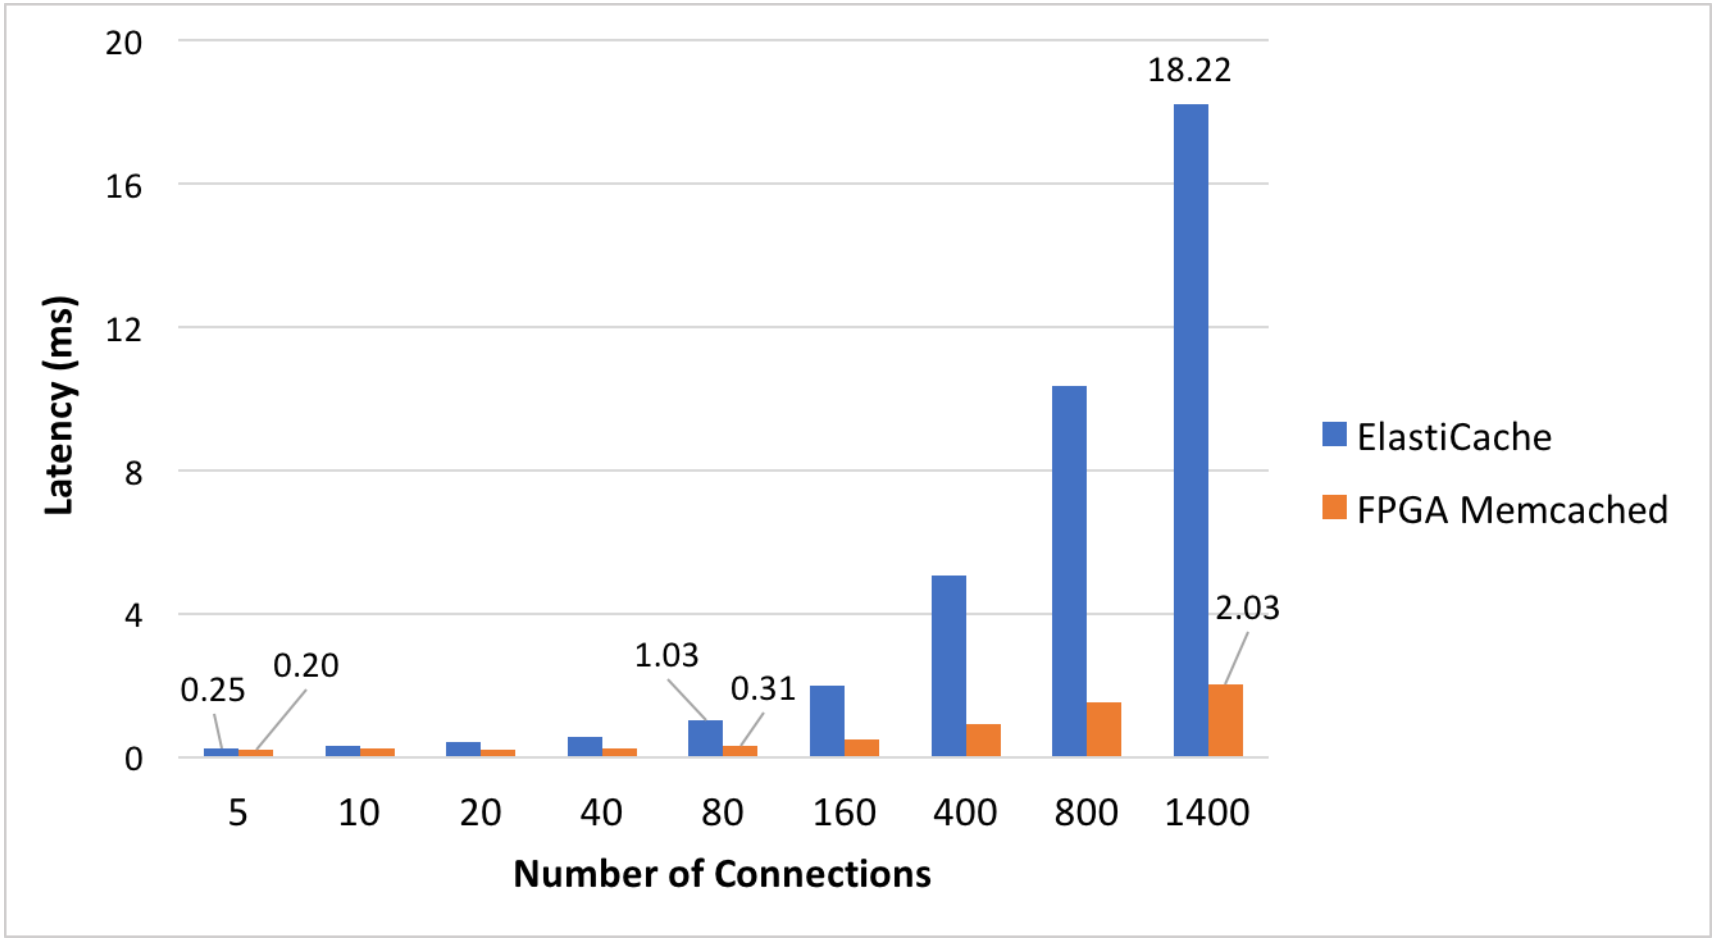

Total Cost of Ownership Results

- Up to **10X** better throughput/$ vs AWS ElastiCache

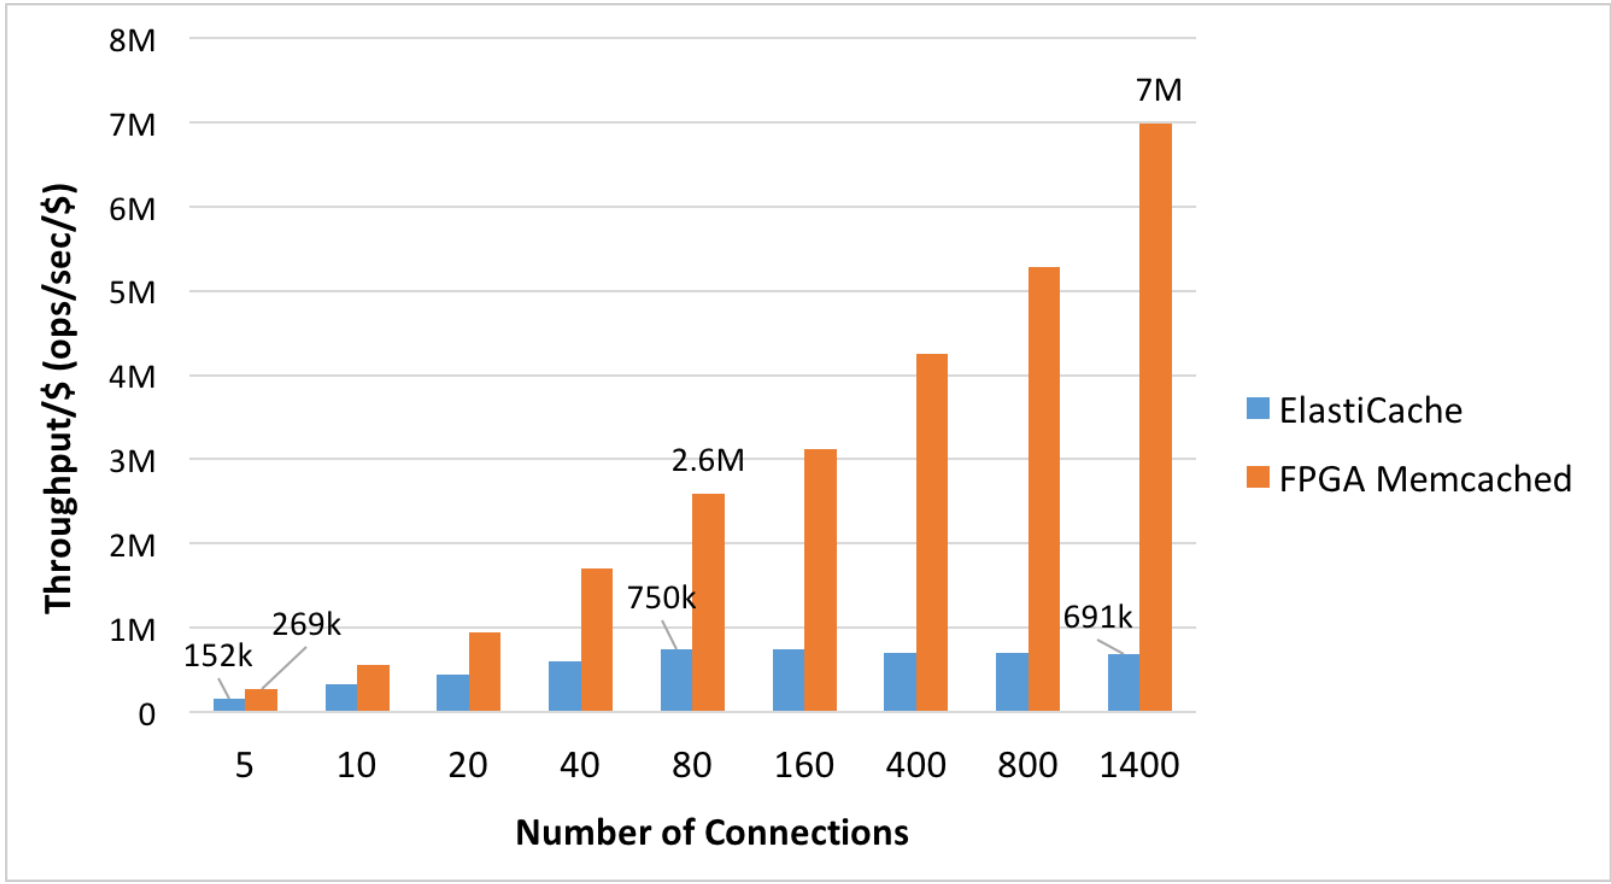

LegUp

Where is the speedup from coming from?

1. We accelerated both TCP/IP network and Memcached completely in FPGA
2. Fully pipelined FPGA hardware – new input every clock cycle
3. Multiple Memcached commands in-flight processed in streaming fashion
4. At high packets/sec, software network stack can become a bottleneck

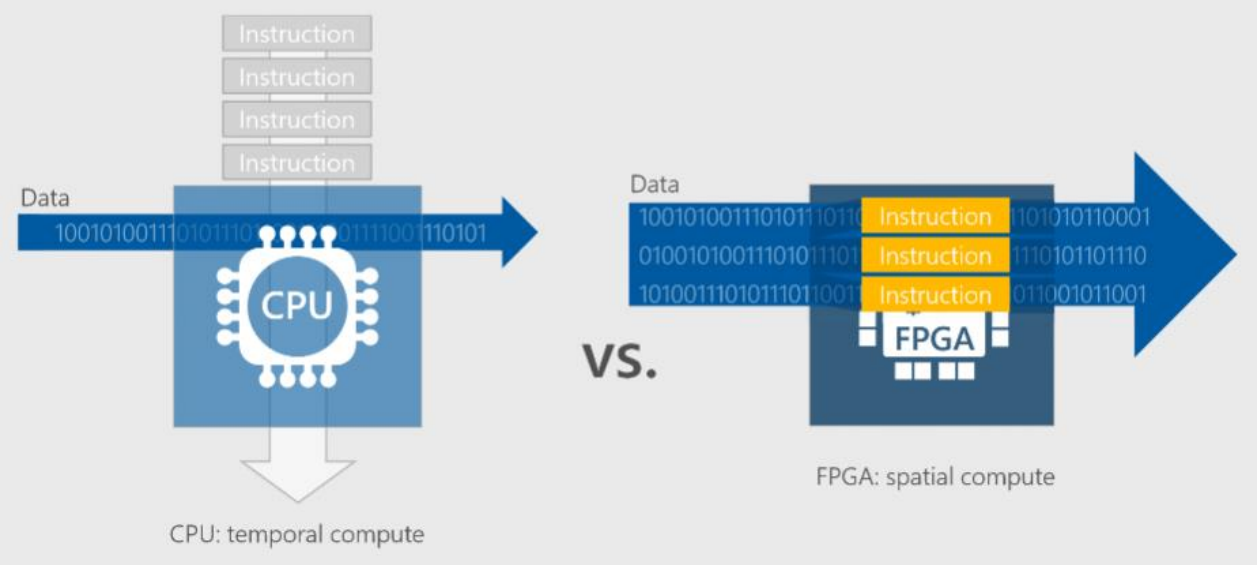

Memcached Demonstration on AWS F1

- Live demo from our website:

  http://www.legupcomputing.com/main/memcached_demo

- Spins up an AWS F1 instance and another client instance

lanny2121 — centos@ip-172-31-33-182:~ — ssh -i ~/aws/Lanny.pem centos@ec2-18-232...

```
F1_Server:~$
```

EC2 Management Console | ElastiCache Management Cons

https://console.aws...

aws | Services | Resource Groups | lian @ 4495-7130-7013

Create | Reboot | Delete | Modify | Manage tags

Filter: Search Clusters... | Viewing 1 of 1 Clusters

Cluster Name	Nodes	Node Type	Zone	Configuration Endpoint
elasticcacheinstance	1 node	cache.r4.4xlarge	us-east-1a	elasticcacheinstance.3pojax.cf

Cluster: elasticcacheinstance

Creation Time: December 6, 2018 at 5:42:53 PM UTC-5

Feedback | English (US) | Privacy Policy | Terms of Use

aws — centos@ip-172-31-37-4:~ — ssh -i James.pem centos@ec2-34-237-197-19.compute-1.amazonaws.com — bash — 157×18

```
MEMTIER_CLIENT:~$
```

lanny2121 — centos@ip-172-31-33-182:~ — ssh -i ~/aws/Lanny.pem centos@ec2-18-232...

```
F1_Server:~$
```

EC2 Management Console | ElastiCache Management Cons

https://console.aws...

aws Services Resource Groups lian @ 4495-7130-7013

Cluster: elasticcacheinstance	**Creation Time:** December 6, 2018 at 5:42:53 PM UTC-5
Configuration Endpoint: elasticcacheinstance.3pojax.cfg.use1.cache.amazonaws.com:11211	**Status:** available
Engine: memcached	**Engine Version Compatibility:** 1.4.34
Node type: cache.r4.4xlarge	**Availability Zones:** us-east-1a
Number of Nodes: 1	**Number of Nodes Pending Creation:**

Feedback English (US) Privacy Policy Terms of Use

aws — centos@ip-172-31-37-4:~ — ssh -i James.pem centos@ec2-34-237-197-19.compute-1.amazonaws.com — bash — 157×18

```
MEMTIER_CLIENT:~$
```

lanny2121 — centos@ip-172-31-33-182:~ — ssh -i ~/aws/Lanny.pem centos@ec2-18-232...

```
F1_Server:~$
```

EC2 Management Console | ElastiCache Management Cons | https://console.aws...

aws | Services | Resource Groups | lian @ 4495-7130-7013

Cluster: elasticcacheinstance

Creation Time: December 6, 2018 at 5:42:53 PM UTC-5

Configuration Endpoint: elasticcacheinstance.3pojax.cfg.use1.cache.amazonaws.com:11211

Status: available

Engine: memcached

Engine Version Compatibility: 1.4.34

Node type: cache.r4.4xlarge

Availability Zones: us-east-1a

Number of Nodes: 1

Number of Nodes Pending Creation:

Feedback | English (US) | Privacy Policy | Terms of Use

aws — centos@ip-172-31-37-4:~ — ssh -i James.pem centos@ec2-34-237-197-19.compute-1.amazonaws.com — bash — 157×18

```
[RUN #1] Preparing benchmark client...
[RUN #1] Launching threads now...
[RUN #1 100%,   8 secs]  0 threads:    10240000 ops, 1245820 (avg: 1237368) ops/sec, 104.44MB/sec (avg: 103.74MB/sec), 13.17 (avg: 13.23) msec latency

32        Threads
32        Connections per thread
10000     Requests per thread


ALL STATS
========================================================================
Type         Ops/sec     Hits/sec   Misses/sec      Latency       KB/sec
------------------------------------------------------------------------
Sets       636437.22          ---          ---     13.23000     88203.36
Gets       636437.22    636437.22         0.00     13.23000     21078.75
Waits           0.00          ---          ---      0.00000          ---
Totals    1272874.45    636437.22         0.00     13.23000    109282.11
MEMTIER_CLIENT:~$
```

lanny2121 — centos@ip-172-31-33-182:~ — ssh -i ~/aws/Lanny.pem centos@ec2-18-232...

```
[F1_Server:~$ start_demo
AFI          0       agfi-002cc4dcf752f2e1b  loaded          0
    ok               0       0x071417d3
AFIDEVICE    0       0x1d0f      0xf000      0000:00:1d.0
EAL: Detected 8 lcore(s)
EAL: No free hugepages reported in hugepages-1048576kB
EAL: Probing VFIO support...
EAL: PCI device 0000:00:03.0 on NUMA socket -1
EAL:   Invalid NUMA socket, default to 0
EAL:   probe driver: 1d0f:ec20 net_ena
EAL: PCI device 0000:00:04.0 on NUMA socket -1
EAL:   Invalid NUMA socket, default to 0
EAL:   probe driver: 1d0f:ec20 net_ena
PMD: eth_ena_dev_init(): Initializing 0:0:4.0
Port 0 MAC: 0e d9 19 5b b3 6c
device_file_name=/dev/edma0_queue_0

Core 0 processing packets. [Ctrl+C to quit]
CPU freq =  2300000000 Hz
Processing on core 0
Time interval = 1.000000s, PPS RX = 0, PPS TX = 0
```

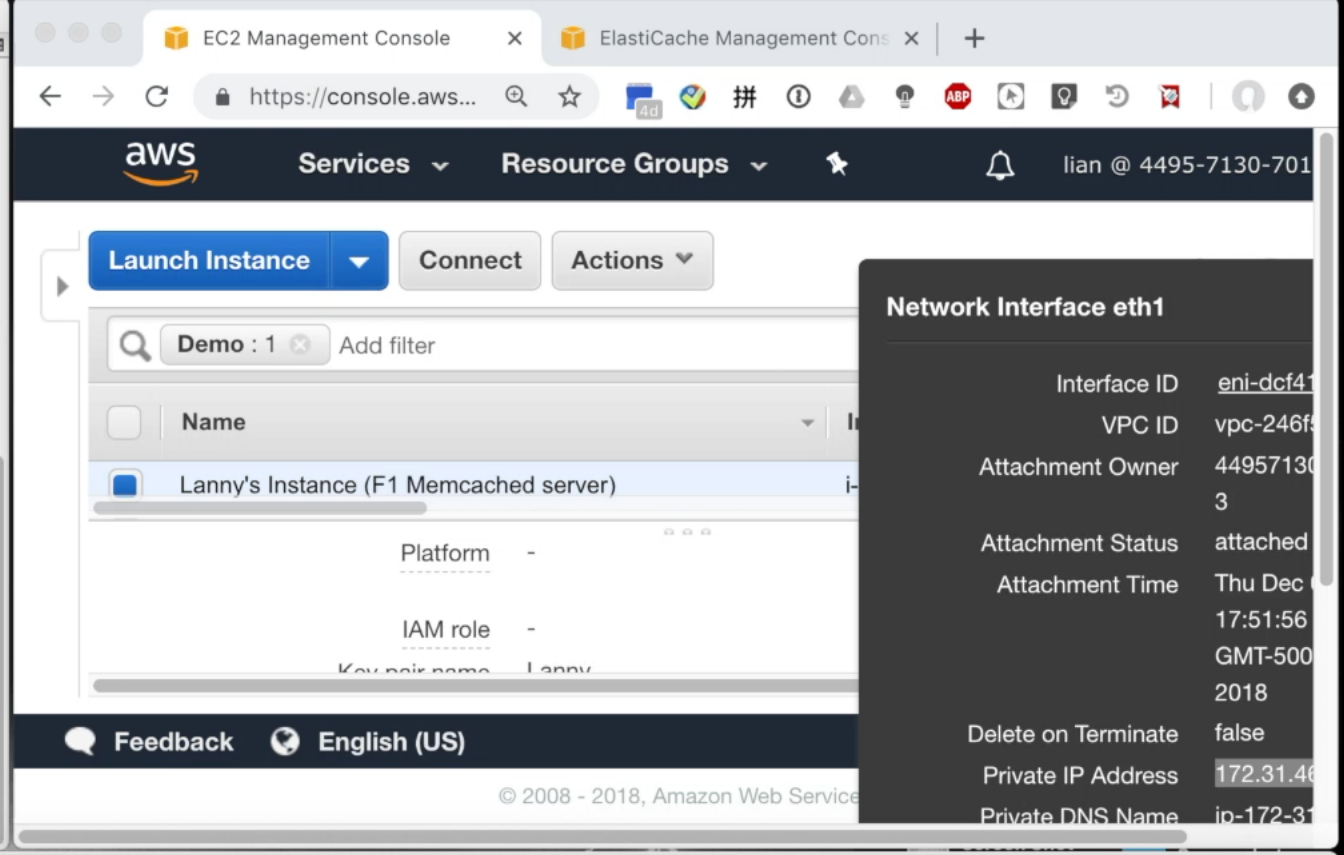

aws — centos@ip-172-31-37-4:~ — ssh -i James.pem centos@ec2-34-237-197-19.compute-1.amazonaws.com — bash — 157×18

```
[RUN #1] Preparing benchmark client...
[RUN #1] Launching threads now...
[RUN #1 100%,   0 secs]  0 threads:    10240000 ops,       0 (avg: 11138885) ops/sec, 0.00KB/sec (avg: 933.91MB/sec),  0.00 (avg:  1.46) msec latency

32        Threads
32        Connections per thread
10000     Requests per thread


ALL STATS
========================================================================
Type         Ops/sec     Hits/sec   Misses/sec      Latency       KB/sec
------------------------------------------------------------------------
Sets      5523383.76          ---          ---      1.46300    765481.66
Gets      5523383.76   5460735.43     62648.33      1.46300    182934.04
Waits           0.00          ---          ---      0.00000          ---
Totals   11046767.53   5460735.43     62648.33      1.46300    948415.70
MEMTIER_CLIENT:~$
```

AWS Cloud-Deployed FPGAs (F1)

- On F1, the FPGA is not directly connected to the network
- CPU is connected to the network and FPGA is connected over PCIe.

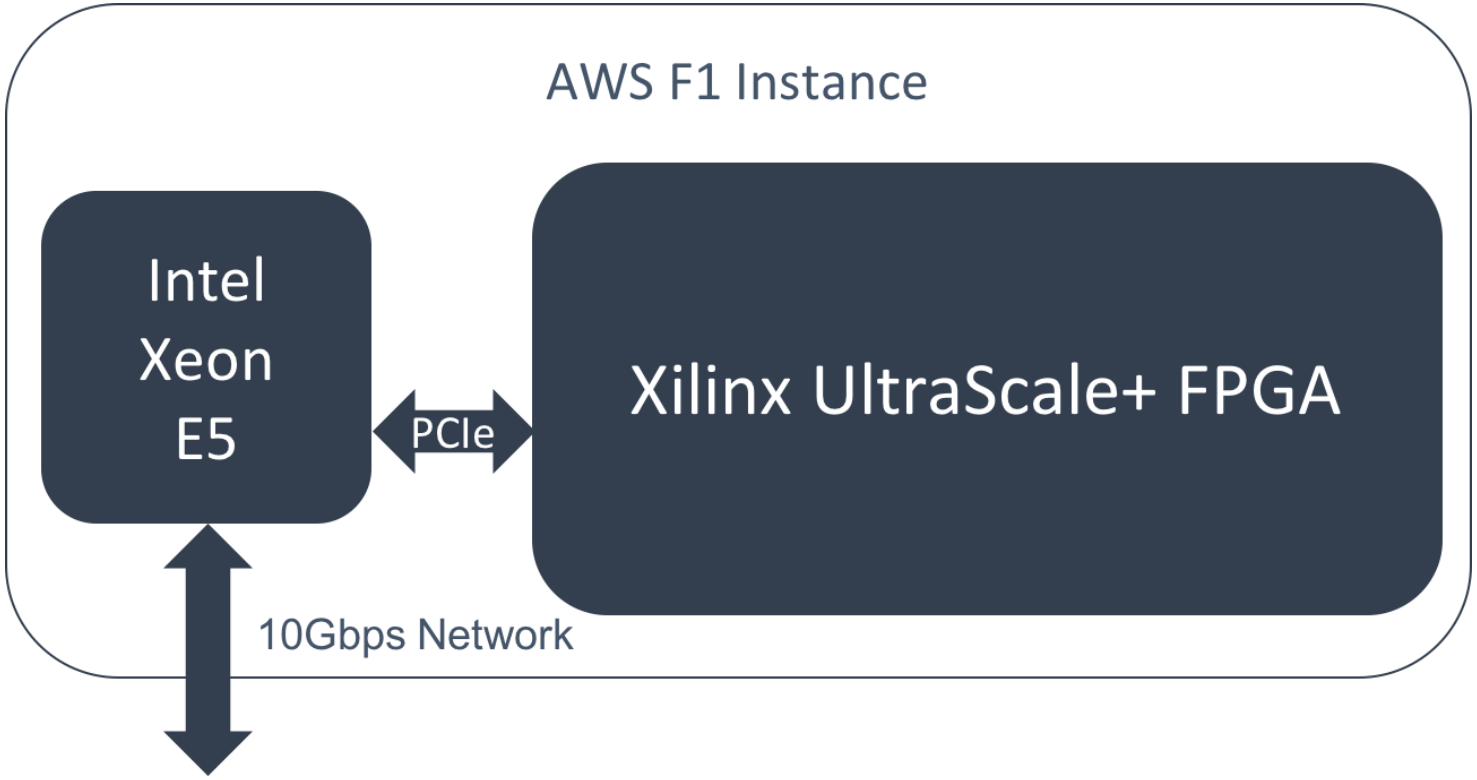

Memcached System Architecture

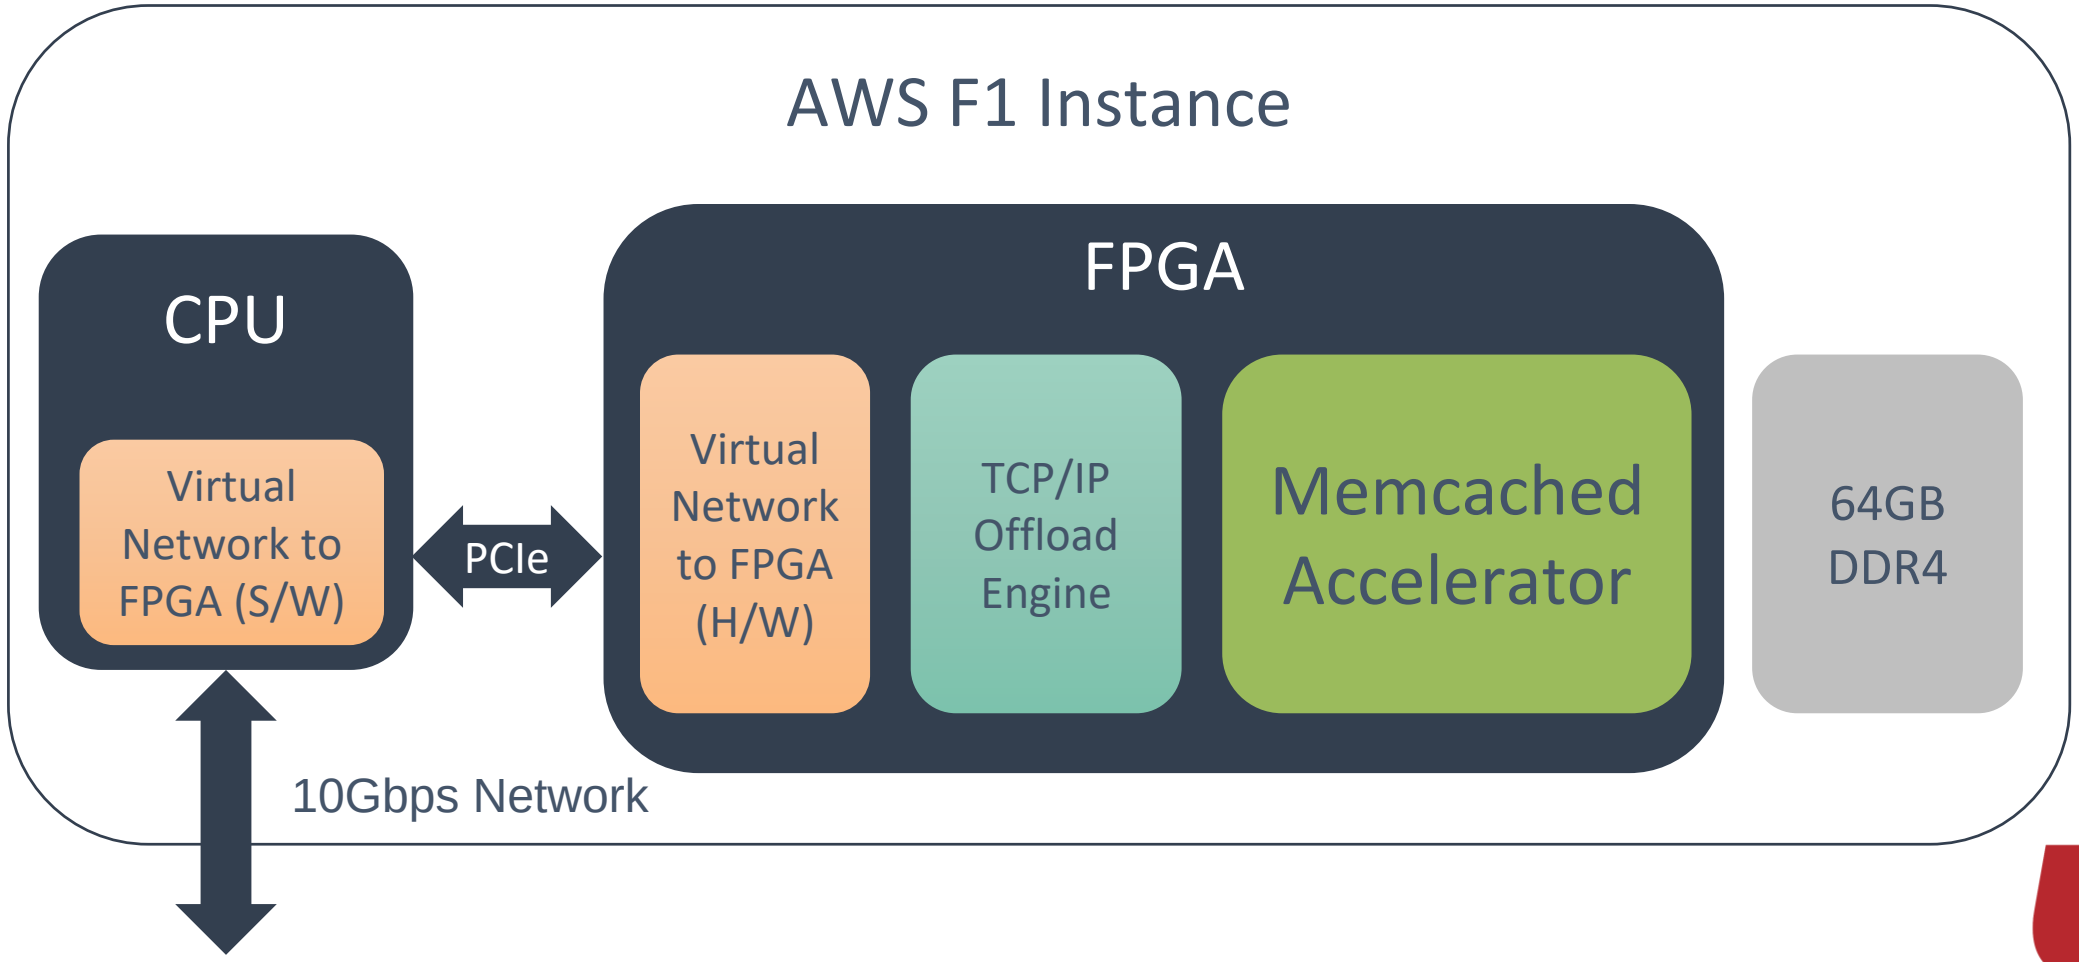

LegUp

Virtual Network to FPGA (VN2F)

- VN2F SW
  - Bypass Linux kernel, send/receive raw network packets, DMA from/to FPGA
- VN2F HW
  - Split/combine DMA data to/from individual network packets
- Each direction takes 20~50us, transfers are overlapped

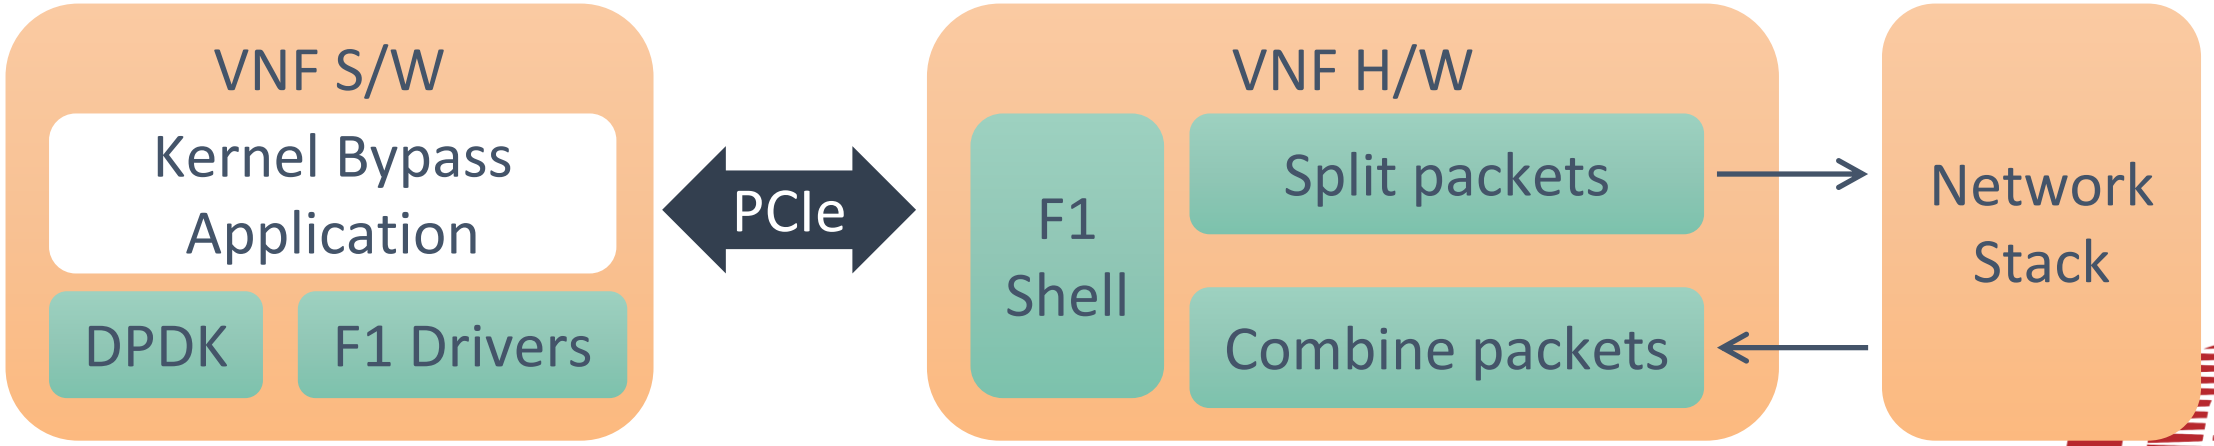

Network Offload: TCP/IP & UDP

- Supports TCP/UDP/IP network protocols
- 10Gbps ethernet support
- 1000s of TCP connections
- Implemented in C++, synthesized by LegUp
- Can be used by other applications
  - Interface with application via AXI-S

* D. Sidler, et al., Scalable 10Gbps TCP/IP Stack Architecture for Reconfigurable Hardware, in FCCM'15

Memcached Core

- The Memcached core is fully pipelined with Initiation Interval of 1
  - Request Decoder block decodes the requests and partitions them into key and value pairs.
  - Hash Lookup hashes keys to hash values and looks up the corresponding addresses
  - Values are stored/retrieved to/from the memory by the Value Store block.
  - Response Encoder creates Memcached responses to return to the clients

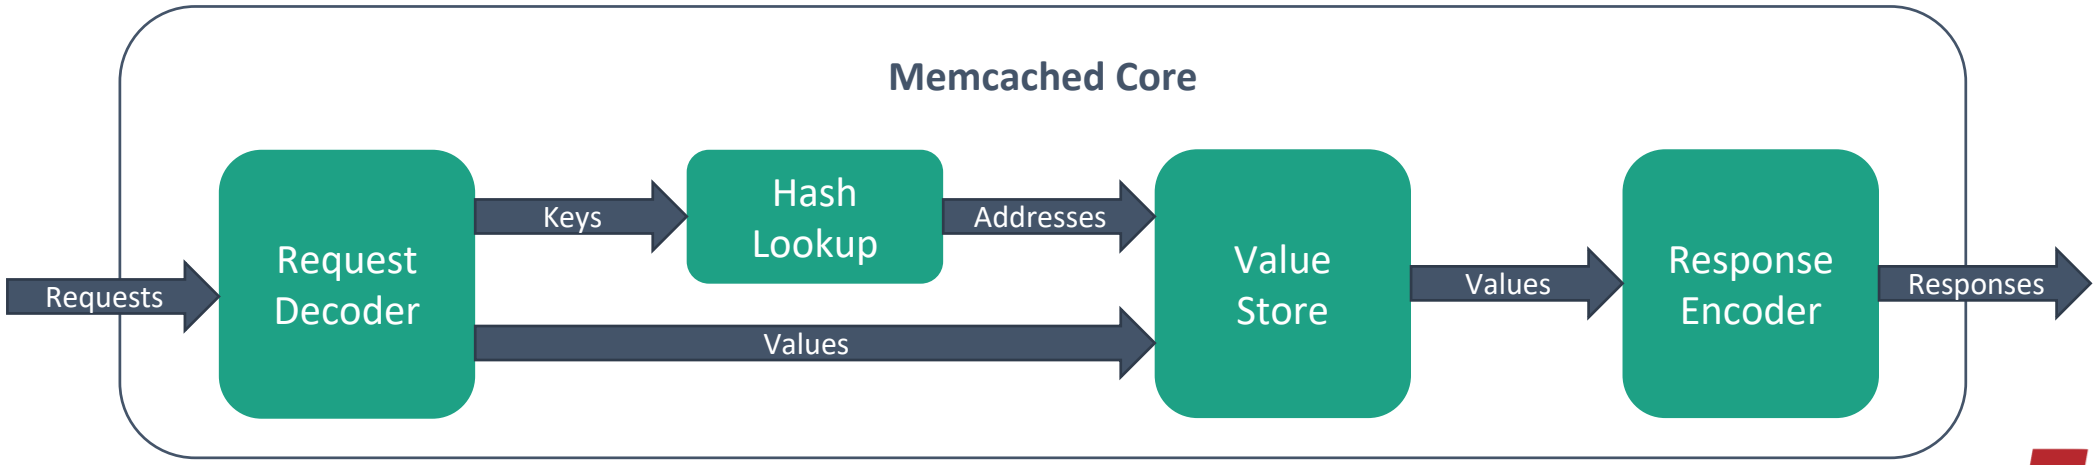

* M. Blott et. al., “Achieving 10Gbps Line-rate Key-value Stores with FPGAs,” in Hot Cloud, 2013

LegUp

Network Bandwidth on AWS F1

- The f1.2xlarge instance has an “Up to 10 Gbps” network
- Bottleneck: bandwidth and PPS

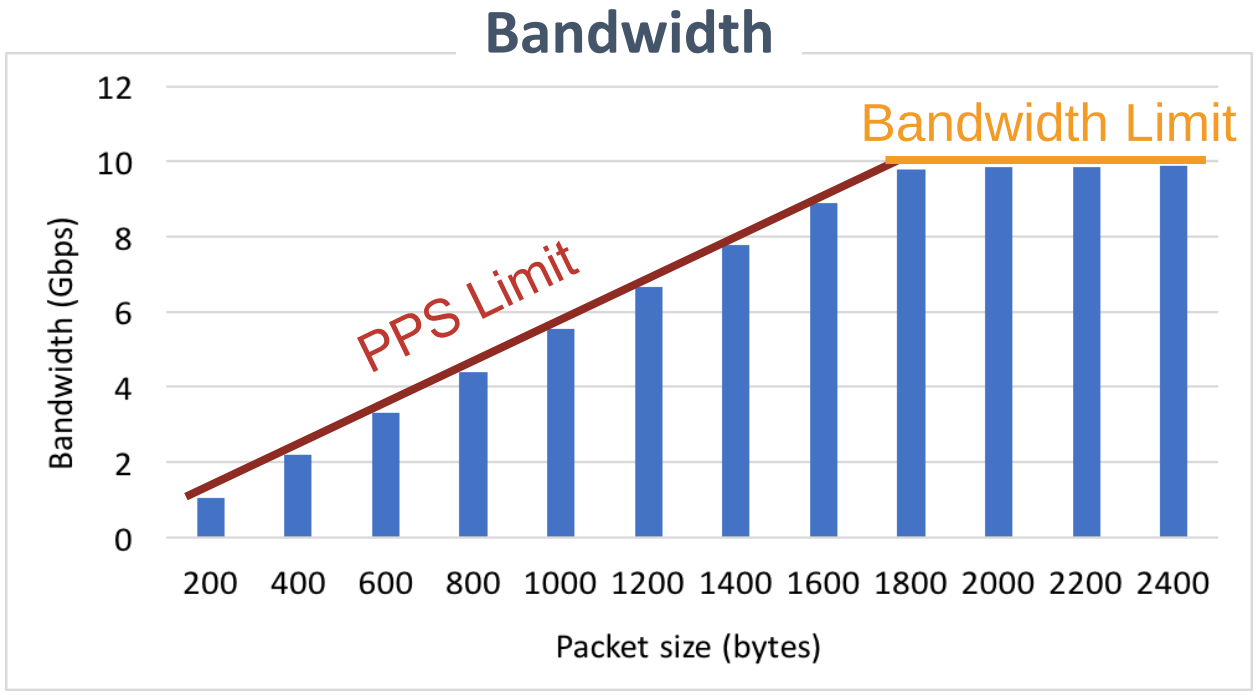

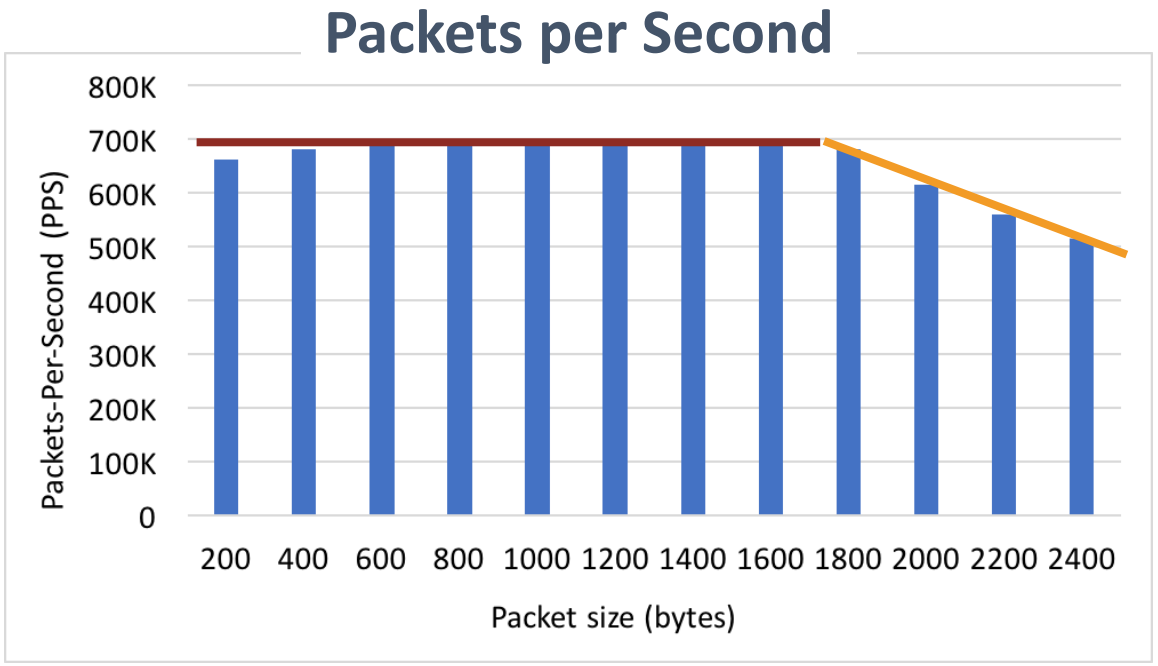

Small packets can’t saturate 10 Gbps network

Max PPS is around 700K

Memcached Request Batching

- Batching in Memcached permits packing multiple requests into a single network packet
  - Reduces packet pro
  - Important feature f neck

Without Pipelining

With Pipelining

Batching is not needed for on-premise FPGAs

- Batching Adapter splits up aggregated requests into individual requests
  - Sends to Memcached core each request in a pipelined fashion

LegUp

Streaming b/w Host & Kernel in SDAccel

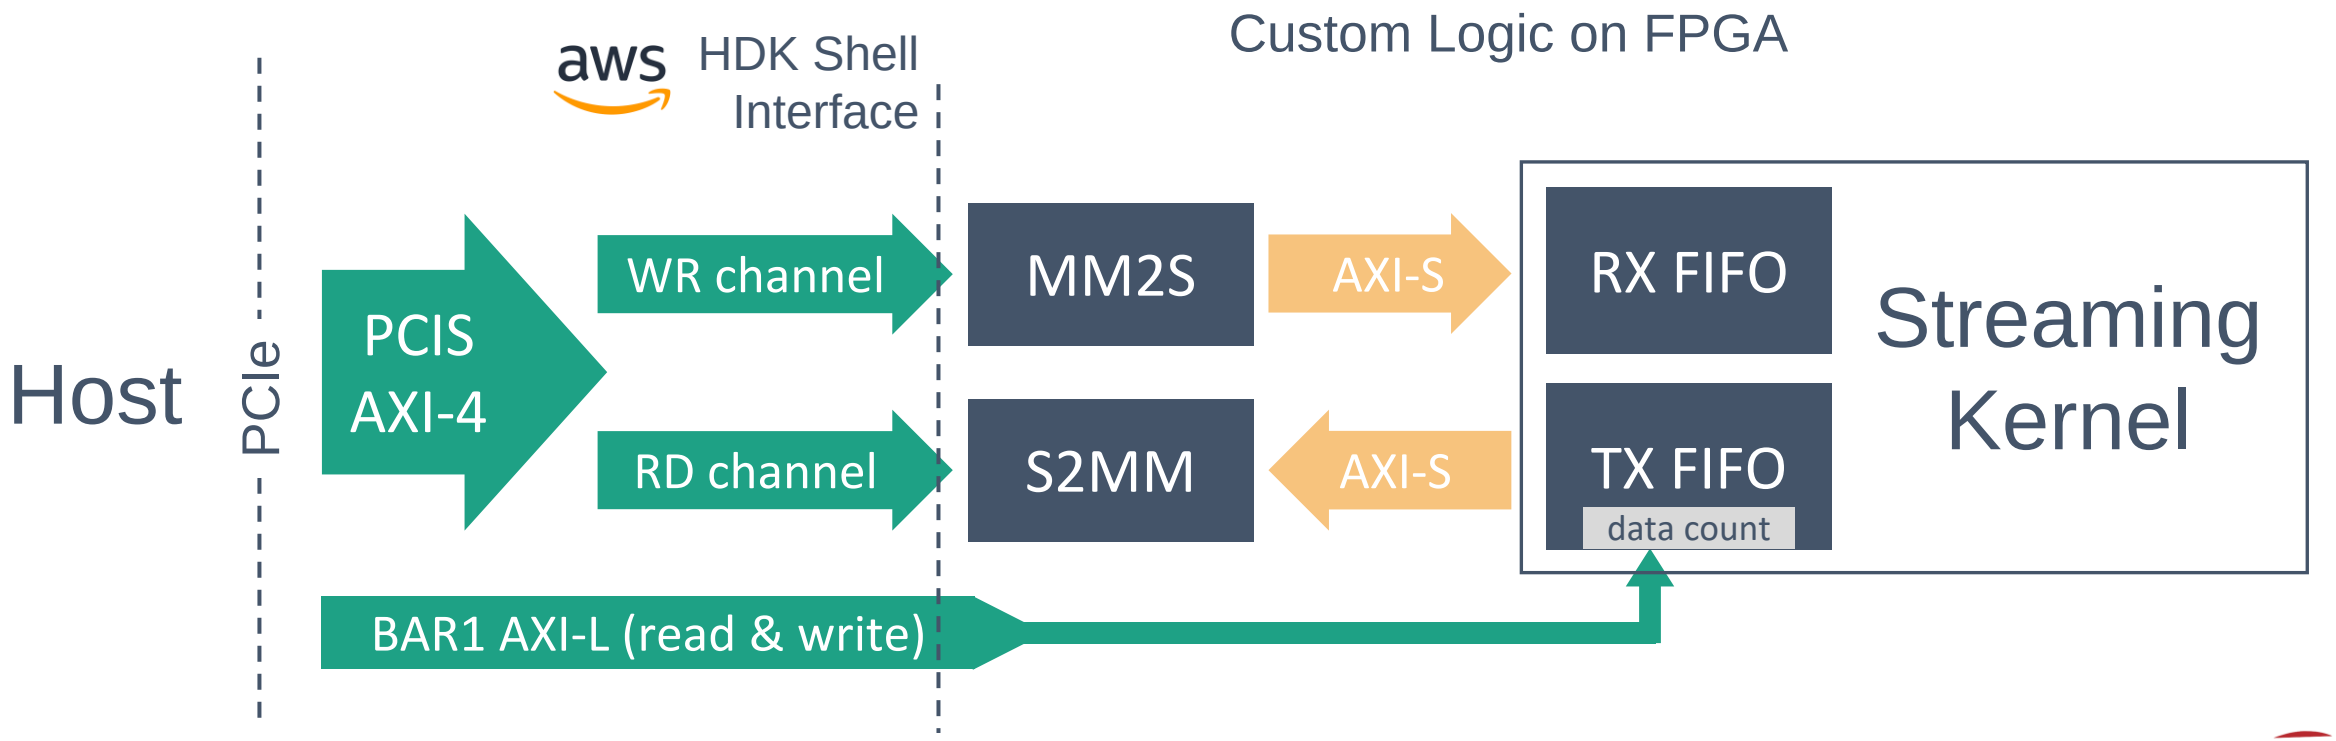

LegUp

Streaming b/w Host & Kernel in SDAccel

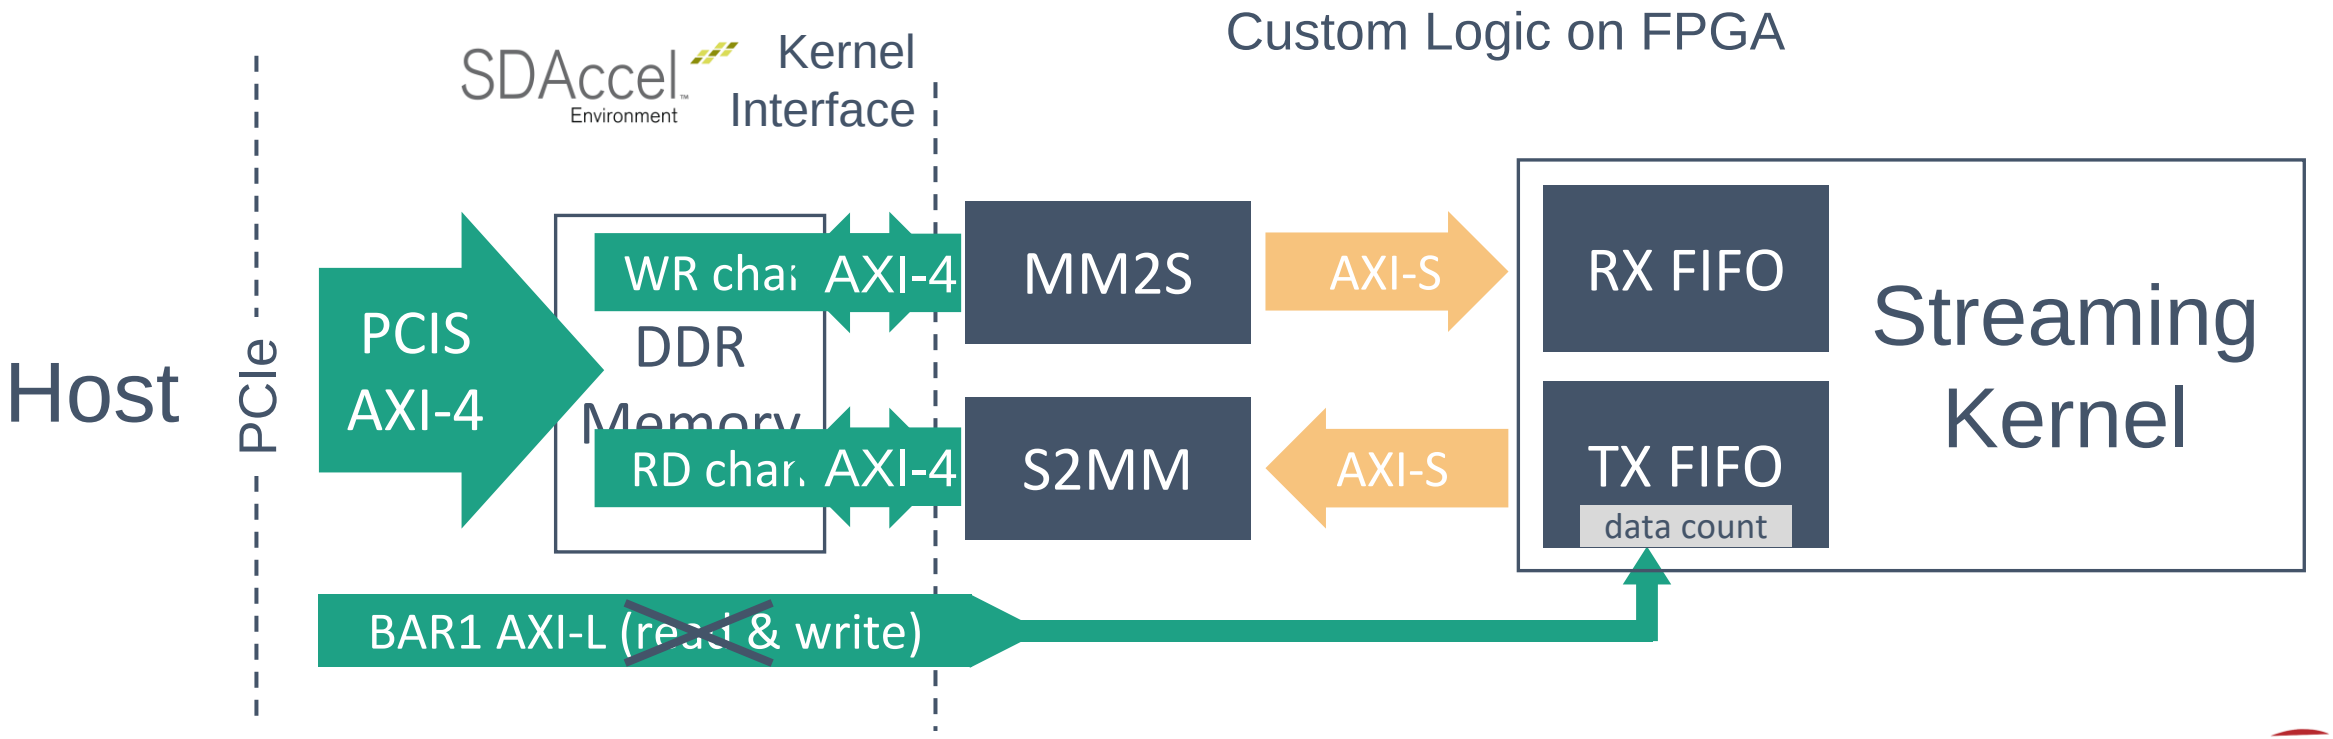

Streaming b/w Host & Kernel in SDAccel

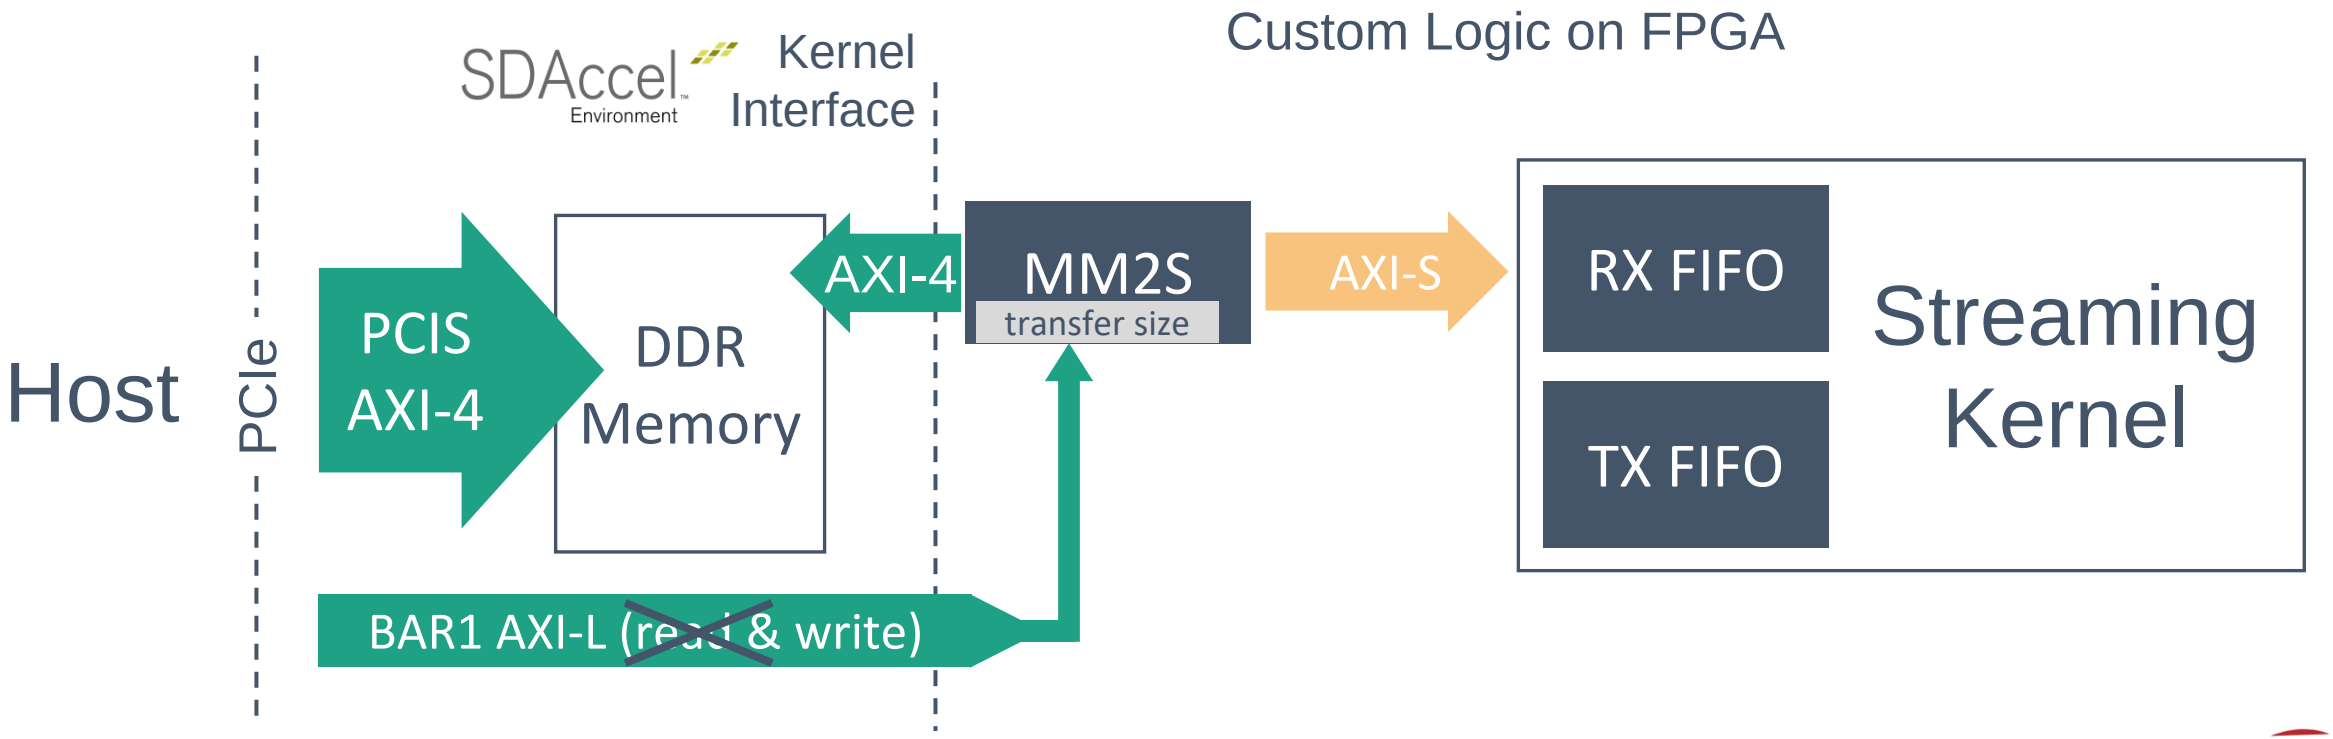

LegUp

Streaming b/w Host & Kernel in SDAccel

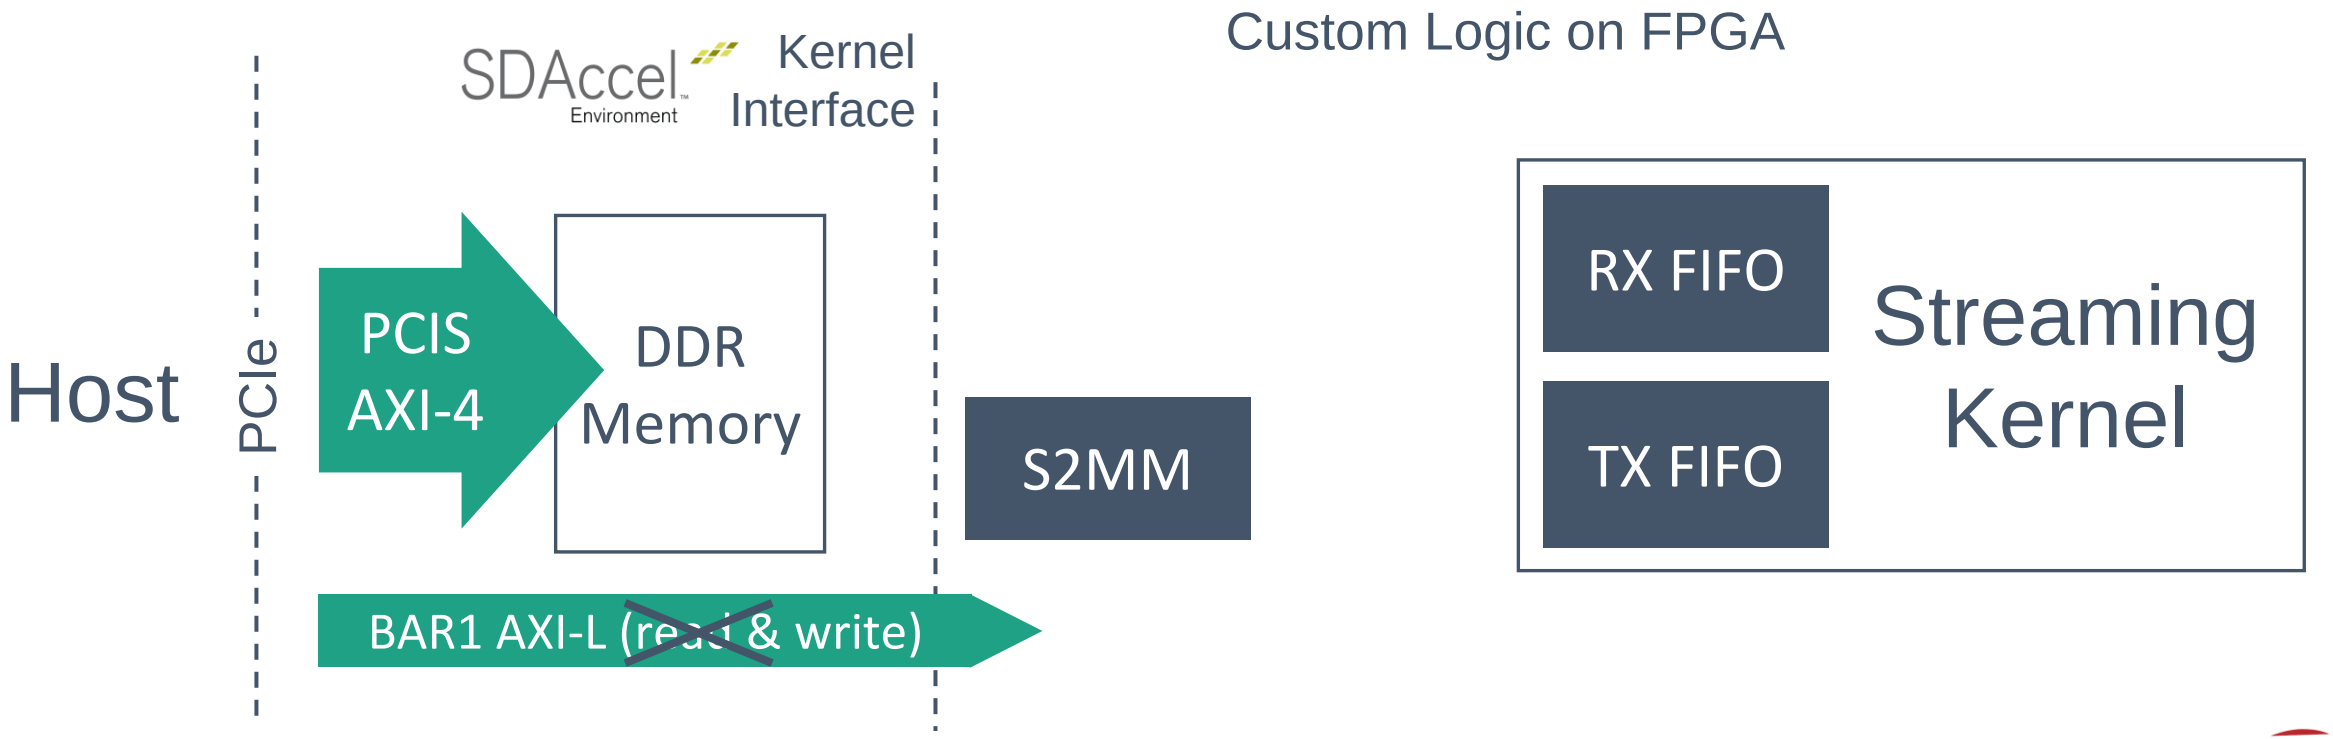

LegUp

Streaming b/w Host & Kernel in SDAccel

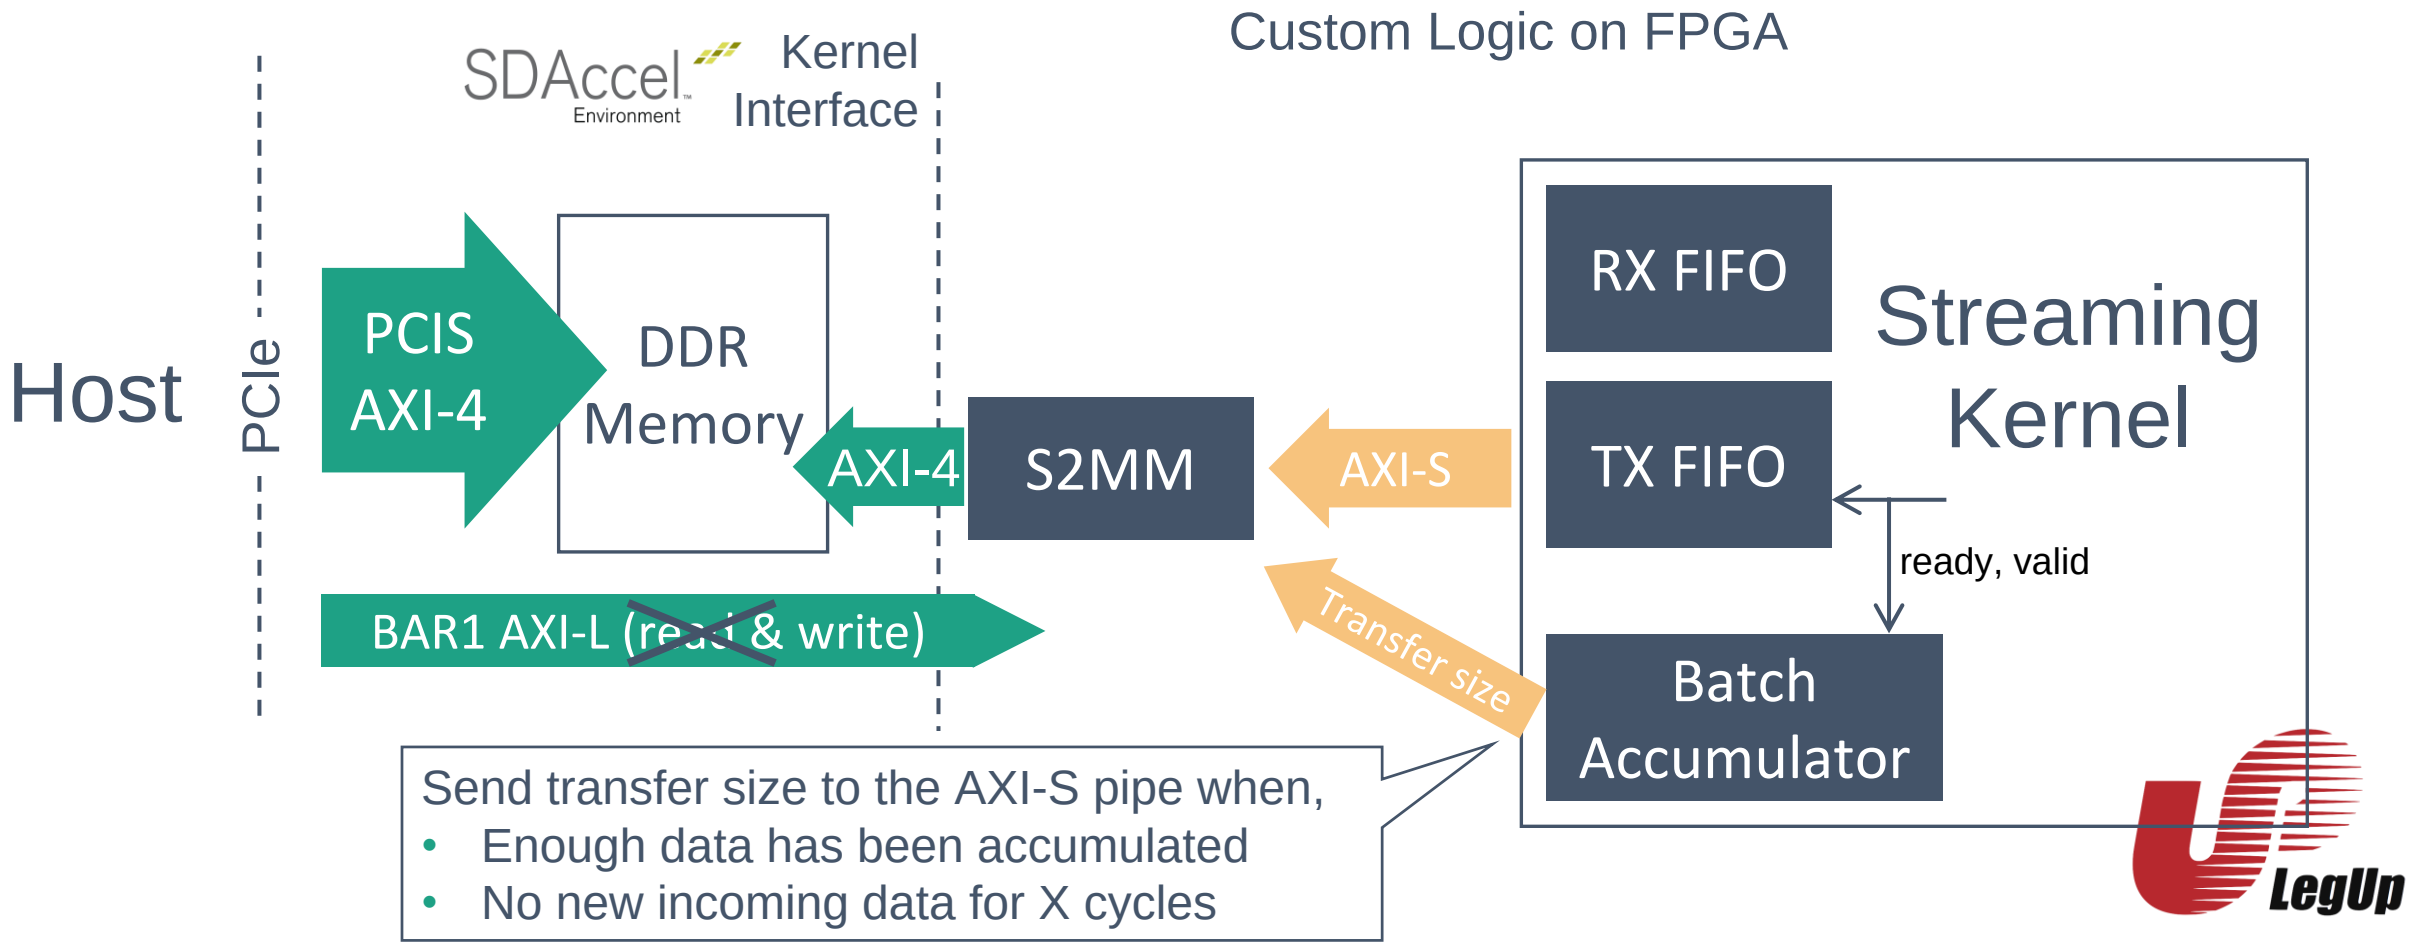

Come talk to us about:

- Memcached Acceleration
- FPGA Network Stack
- SDAccel Streaming Handler
- LegUp high-level synthesis tool
- Any other FPGA acceleration needs

Andrew Canis & Ruolong Lian

www.LegUpComputing.com

info@legupcomputing.com | 647-834-6654 | Toronto, Canada